

BEST AVAILABLE COPY**Disclosure SJO8-1999-0443**

Created By: Kevin Smith **Created On:** 12/13/99 12:59:01 PM
Last Modified By: Marie Turley **Last Modified On:** 02/14/2000 02:39:07 PM

*** IBM Confidential ***

Required fields are marked with the asterisk (*) and must be filled in to complete the form.

Summary

Status	Search Results Received
Processing Location	SJO
Functional Area	Test & Performance Systems - Truskowski
Attorney/Patent Professional	Randall J Bluestone/San Jose/IBM
IDT Team	Clod Barrera/San Jose/IBM; Brian Smith/San Jose/IBM; Wynn Price/San Jose/IBM; Homayoun Samadi/San Jose/Contr/IBM
Submitted Date	12/13/99 03:43:08 PM
Owning Division	SSD
PVT Score	To calculate a PVT score, use the 'Calculate PVT' button.
Lab	
Technology Code	

Inventors with Lotus Notes IDs

Inventors: Kevin Smith/San Jose/IBM

Inventor Name > denotes primary contact	Inventor Serial	Div/Dept	Manager Serial	Manager Name
Smith, Kevin E. >	155051	354435A	323849	Smith, Bill (Brian)

Inventors without Lotus Notes IDs**IDT Selection**

IDT Team:	Attorney/Patent Professional:
Cloid Barrera/San Jose/IBM Brian Smith/San Jose/IBM Wynn Price/San Jose/IBM Homayoun Samadi/San Jose/Contr/IBM	Randall J Bluestone/San Jose/IBM

Response date: 01/13/2000

Main Idea**Title of disclosure (in English)**

Efficient mixing of sequential prefetches with random access data in a preexisting LRU cache

Idea of disclosure

1. Describe your Invention, stating the problem solved (if appropriate), and indicating the advantages of using the Invention.

Given an existing LRU cache for which it is impossible, or not desirable, to alter it in such a way as to efficiently accommodate sequential-oriented prefetched data this invention describes a process for prefetching data in such a way as to work efficiently with random accesses in said LRU cache.

Efficient mixing of sequential prefetches with random access data in a preexisting LRU cache - continued

The advantage of this technique is that it is not always possible, or desirable, to alter the replacement algorithm of an existing LRU cache. This invention is designed to work with a preexisting LRU cache without alteration. This applies even when the LRU cache is across a physical boundary, such as a computer host attempting to make use of the cache within an I/O subsystem for efficient caching of prefetched data.

2. How does the invention solve the problem or achieve an advantage, (a description of "the invention", including figures inline as appropriate)?

The invention describes a process for estimating the single reference residency time of elements within a cache and then using that information to model the utility of keeping prefetched data for I/O streams which are expected to benefit from cache prestaging.

The I/Os eligible for signalling prefetching are recorded in a computer model. Because of the single reference residency prediction the number of simulated prefetches remains relatively small. A sequential prefetch "count" is recorded with each element in the computer model, providing a basis for decision with regards to invoking an actual prefetch into the LRU cache in question. A measure of "goodness" (based on this "count") for prefetching exists such that the threshold value for prefetching eligibility can increase or decline depending upon the LRU cache's hit ratio performance.

3. If the same advantage or problem has been identified by others (inside/outside IBM), how have those others solved it and does your solution differ and why is it better?

I am not aware of prefetching algorithms which do not require some degree of modification to the LRU cache being prefetched into. A common "solution" is to maintain a sequential "count" (as described above) with each cache element, thus providing a measure of "goodness" for prefetch eligibility, however in every case that I am familiar with the "count" information needs to be added to the cache element information ^{in cache, not rewrite}, comprising the LRU cache, thus requiring a modification to the LRU cache in question.

4. If the invention is implemented in a product or prototype, include technical details, purpose, disclosure details to others and the date of that implementation.

This invention is not currently implemented in any product. I have an internal prototype which implements the algorithm. This internal prototype which is not associated with any specific product was created on 12/10/1999.

*Critical Questions (Questions 1 - 7 must be answered)

Question 1

On what date was the invention workable? 12/10/99. Please format the date as MM/DD/YYYY.
(Workable means, i.e., when you know that your design will solve the problem)

Question 2

Is there any planned or actual publication or disclosure of your invention to anyone outside IBM? Yes No

If yes, Enter the name of each publication or patent and the date published below:

Publication/Patent:

Date Published or Issued:

Are you aware of any publications, products or patents that relate to this invention?

Yes

No

If yes, Enter the name of each publication or patent and the date published below:

Publication/Patent:

Date Published or Issued:

Question 3

Has the subject matter of the invention or a product incorporating the invention been sold, used internally in manufacturing, announced for sale, or disclosed in a proposal?

Yes

No

Efficient mixing of sequential prefetches with random access data in a preexisting LRU cache - continued

Is a sale, lease, or manufacturing product announcement or proposal planned?

Yes No

If yes, identify the product if known and indicate the date or planned date of sale, announcement, or proposal and to whom the sale, announcement, or proposal has been or will be made.

Product: _____
Version/Release: _____
Confidentiality: _____
Date: _____
To Whom: _____

If more than one use cut and paste and append as necessary in the field provided.

Question 4

Was the subject matter of your invention or a product incorporating your invention used in public, e.g., outside IBM or in the presence of non-IBMers?

Yes No

If yes, give a date. Please format the date as MM/DD/YYYY

Question 4

Have you ever discussed your invention with others not employed at IBM?

Yes No

If yes, identify individuals and date discussed. Fill in the text area with the following information: the names of the individuals, the employer, date discussed, under CDA, and CDAM.

Question 5

Was the invention, in any way, started or developed under a government contract or project?

Yes No Not Sure

If Yes, enter the contract number.

Question 7

Was the invention made in the course of any alliance, joint development or other contract activities?

Yes No Not Sure

If Yes, enter the following. Name of Alliance, Contractor or Joint Developer
Contract ID number _____
Relationship/Contract name _____
Relationship contact E-mail _____
Relationship contact phone _____

Question 9

Have you submitted, or are you aware of, any related disclosure submission?

Yes No

If Yes, please provide the title and docket or disclosure number below.

Efficient mixing of sequential prefetches with random access data in a preexisting LRU cache - continued

Question 9:

What type of companies do you expect to compete with inventions of this type? Check all that apply.

- Manufacturers of enterprise servers
- Manufacturers of entry servers
- Manufacturers of workstations
- Manufacturers of PCs
- Non-computer manufacturers
- Developers of operating systems
- Developers of networking software
- Developers of application software
- Integrated solution providers
- Service providers
- Other (Please specify below)

Patent Value Tool (Optional - this may be used by the Inventor and attorney to assist with the evaluation of your invention.)

(The Patent Value tool can be used by you or the evaluation team to determine the potential licensing value of your invention.)

The Patent Value Tool has not yet been used to calculate a score.

Evaluation

This evaluation was entered by Sarah Hamel/San Jose/Contr/IBM on 01/14/2000

Team Evaluation

What is the team's evaluation of this disclosure? Search

Date rated: 01/14/2000

Evaluation Comments

Search Information

Date sent: 01/14/2000	Target completion date:	Actual completion date: 02/09/2000
Who was the search sent to? (This area is to designate a Local Searcher name or WAIFL) JWM		
Send search request to:	Search Type:	
Patent Office/Attorney/IBM/IBMUUS	<input checked="" type="checkbox"/> Patentability <input type="checkbox"/> Clearance <input type="checkbox"/> Validity <input type="checkbox"/> State of Art	
How far back in time will the search cover?		
Features to be searched: Please see Disclosure. *PLEASE NOTE: ADDITIONAL INFORMATION (BIFOLDS) FOLLOWING THIS PAGE WILL BE INCLUDED*		

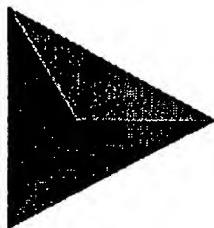
Post Disclosure Text & Drawings

Enter any additional information relating to this disclosure below:

(Form Revised 12/17/97)

Efficient Mixing of Sequential Prefetches With Random Access Data in a Preexisting LRU Cache

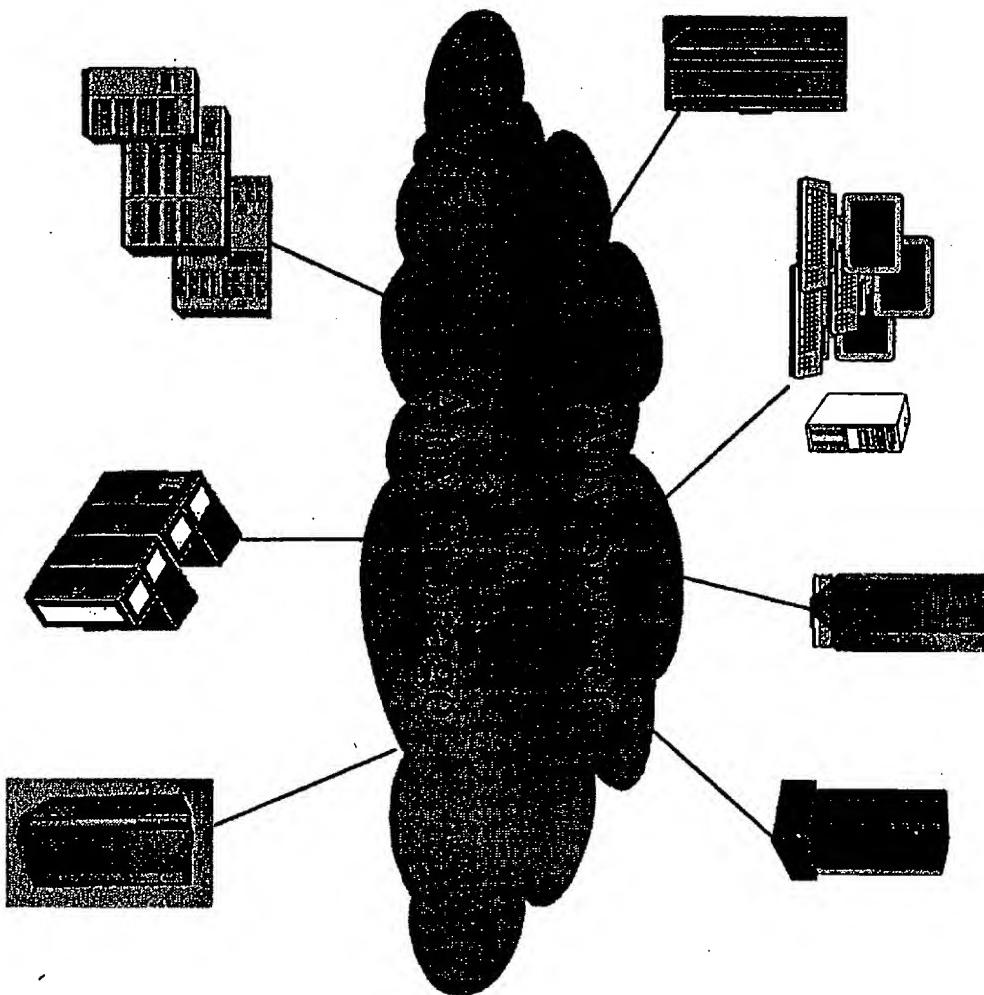
Patent Review Board
IBM Confidential



IBM Confidential**Efficient Prefetching w/LRU...**

- Title:
 - ▶ ■ Product Name:
 - First Date
 - Disclosed:
- Efficient Mixing of Sequential Prefetches with Random Access Data in a Preexisting LRU Cache.
- Problem Solved:
 - ▶ ■ Prefetch optimization in a black-box LRU caching system.

Example use of the process in a SAN environment



IBM Confidential

Problem

- LRU cache logic sometimes inaccessible or undesirable to alter.
- LRU cache often accessed by external hosts or multiple hosts.
- Prefetched data is ideal for sequential accesses, which works poorly with LRU.
- Identifying prefetch candidates difficult.

Solution

- Solution is to:
 - Use model of prefetching effects to make decisions.
 - Estimate single ref. residency time.

IBM Confidential

Efficient Prefetching w/LRU

- Benefit over prior art.
- Competitors:
 - ▶ Improved identification of prefetch candidates.
 - Modification of LRU not required.
- Will others want to use it?
 - Consider a SAN with a cached control unit to preload into.
- Discoverable?
 - ▶ I/O requests in addition to workload together with a moving prefetch criteria.
- Alternatives?
 - Static prefetch determination.

IBM Confidential

The process for efficient prefetching w/LRU

- Determine cache size.
- Periodically fetch Hit Ratio, cached I/O rate and estimate the SRRT.
- For each I/O check the model's buffer, use previous element's sequential count.
- If above a dynamic threshold, prefetch.

The process for efficient prefetching w/LRU (continued)

- Load I/O requests + prefetches into model's buffer. Update LRU position on hit. Discard on overflow.
- Keep track of overflow point for each of multiple prefetch threshold criteria, counting prefetched "hits" for each criteria.
- If alternate prefetch criteria value is better, then consider adjustment.

IBM Confidential

Description of Invention

- Inputs:
 - Cache Size
 - periodic:
 - Hit Ratio feedback
 - Cached I/O Rate feedback
- Output:
 - Prefetch candidates identified
 - Optimal prefetch threshold identified
 - Value of prefetching quantifiable

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.